

CORRELATING NETWORK INFORMATION AND INTRUSION INFORMATION TO FIND THE ENTRY POINT OF AN ATTACK UPON A PROTECTED COMPUTER

FIELD OF THE INVENTION

The present invention is related to the field of data processing security, and more particularly to a
5 method for determining the entry point or exit point of an attack by a vandal upon a device such
as a computer or a web server that is protected by an intrusion detection system.

BACKGROUND

Computer activities are sometimes subject to electronic vandalism. For example, a vandal or
hacker may attack an Internet web server by flooding it with a torrential flow of disruptive
10 messages that overload the server to the point of functional failure. Attacks of this kind are
called “denial of service” attacks. During a denial of service attack, the vandal may fraudulently
assume a number of different electronic identities, often by including messages in the disruptive
flow that have a variety of source addresses. To combat such attacks, a web server may rely
upon intrusion detection equipment that examines incoming messages. Such equipment detects
15 the onslaught of a vandal’s attack, reads source addresses including the addresses that the
attacker usurps and fraudulently re-uses, and issues alerts upon detection of all messages that
exhibit characteristics of the attack.

Computers are also subject to other kinds of attacks, for example attacks that are mounted by parties commonly known as hackers. A vandal such as a hacker may attempt to intrude upon a computer in order to steal information in an act of industrial espionage, or to implant a virus, or to alter records to the detriment or benefit of someone's interests or reputation. Again, to combat such activities, computers may be monitored and protected by intrusion detection systems.

Intrusion detection systems are effective and useful for their intended purposes. Unfortunately, however, the protected computer, the intrusion detection system, and the associated network equipment, such as firewalls and routers, are separate devices today whose operations are essentially uncoordinated. This lack of coordination limits the capability of the combined system to respond to attacks in any useful way except for detecting the presence of an attack and attempting to limit – by filtering – any damage that might result.

For example, when an intrusion detection system detects a denial of service attack upon a computer, the attack may be blocked by a firewall to the extent possible by source-address filtering. Blocking such an attack may, however, have adverse and unintended consequences. In particular, the use of protective filtering may play into the hands of a vandal who resorts to “spoofing.” A spoofer is an attacker who uses a fake source address that fraudulently identifies the spoofer as another source. Spoofing attacks may have serious consequences, for example when the spoofer usurps the source address of a web server's most important customer. In such instances, the administrators of the network of the targeted web server may unknowingly decide to filter-out all messages that bear the customer's source address, inadvertently including

messages actually sent by the customer, using protective equipment such as firewalls and routers.

Consequently, the web server experiences both the trauma of an attack and the adverse consequences that come with mounting a defense that filters-out legitimate messages sent by the server's most important customer. Moreover, Denial-of-Service (DoS) attacks are generally characterized by extraordinary volumes of data sent to the targeted system or network.

Consequently, even if a firewall were configured to block the DoS traffic, the firewall would be quickly overwhelmed.

As the attack wears on, technicians pore over volumes of data dumped by the intrusion detection system and the network equipment, in an *ad hoc* attempt to determine the attack's entry point into the protected device. Unfortunately, with such an unstructured approach the entry point may not be found for several hours, during which time the intrusion detection system's filtering impedes the operation of the computer or web site under attack.

In other situations, the attack itself may be short lived, for example as in a hacker's attempt to steal information from a protected computer. By the time technicians have completed their unstructured situational analysis, the attack may be over and the hacker may have succeeded before his efforts could be blocked by filtering the computer's outbound flow of information.

Thus there is a need for an improved defense against vandals that provides a quick and efficient way of determining the entry point of an attack upon a device that is protected by an intrusion detection system, so that measures may be taken to stop the attack closer to its source or to aid in

forensic investigations that follow.

SUMMARY OF THE INVENTION

The present invention provides a way of determining the entry point of an attack by a vandal such as a hacker upon a protected device such as a computer or a server such as a web server that operates under the protection of an intrusion detection system. According to the present invention, intrusion detection information regarding an attack and network information regarding the attack are correlated, and from this information the entry point of the attack into the protected device is identified. Thus, the entry point or exit point of the DoS traffic may be identified as far back in the network as possible – i.e., as far as possible away from the protected device.

One embodiment of the invention pertains to a web server that is protected by an intrusion detection system and connected to the Internet through a router. When the intrusion detection system detects an attack upon the web server, the intrusion detection system provides intrusion information to a correlation engine, which information may include the source address, destination address, protocol type of messages included in the attack, and so forth. The correlation engine receives network information from the router, which information may be a router table that includes the source address, destination address, protocol type, logical input port identifier, and logical output port identifier of each connection made through the router.

The correlation engine then correlates the intrusion information and the network information to deduce the logical entry point of the attack. For example, the correlation engine may correlate the intrusion information with the network information by finding, in the router table, occurrences of the intrusion information's source address, destination address, or protocol type.

5 When a match is found between a record of the router table and the intrusion information, the correlation engine concludes that the logical entry point of the attack is the logical input port and logical output port that support the connection through the router. By consulting a mapping table, the correlation engine then identifies physical entry point of the attack by identifying the router's physical ports that are associated with the logical ports of the entry point. The mapping
10 table may be included in the router table, or may be stored separately in the router or elsewhere. An alert identifying the entry point of the attack is sent to network management center, which may then block the attacking traffic from reaching the affected ports, or which may work with service providers to stop the attack closer to its source.

Thus the present invention provides an effective and efficient way of identifying the entry point
15 of an attack upon a device such as a computer or a web server that is protected by an intrusion detection system. The present invention also encompasses situations involving a plurality of routers that are connected in a web or a mesh, and the methods of the present invention are used to identify the entry points of the attacking traffic on all the affected routers. These and other aspects of the invention will be more fully appreciated when considered in light of the following
20 detailed description and drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1, which is a block diagram illustrating exemplary use of the present invention, includes a protected device such as a web server, network equipment such as a router, an intrusion detection system such as a network based intrusion detection system, and a correlation engine.

5 FIG. 2 shows an exemplary structure of intrusion information provided by the intrusion detection system of FIG. 1.

FIG. 3 shows an exemplary structure of network information provided by the network equipment of FIG 1.

10 FIG. 4 shows an exemplary structure of a mapping table that maps logical ports to physical ports of the network equipment of FIG. 1.

FIG. 5 is a flowchart showing aspects of the operation of the invention.

DETAILED DESCRIPTION OF THE INVENTION

The present invention provides a way of determining the entry point of an attack upon a device that is protected by an intrusion detection system. A correlation engine correlates network

information and intrusion information, and thereby deduces the logical entry point of the attack, and from the logical entry point of the attack deduces the physical entry point of the attack.

FIG. 1 shows an intrusion detection system 130 that protects a protected device 100 against deliberate attacks by vandals or inadvertent attacks by others having the apparent characteristics of a vandal's attack. Here, the protected device 100 may be a computer, a server such as a web server, or other similar devices. Although the intrusion detection system 130 of FIG. 1 shows the use of network-based intrusion detection equipment for the purpose of ready discussion, the present invention encompasses the use of other kinds of intrusion detection equipment as well, including host-based intrusion detection equipment, application-based intrusion detection equipment, and so forth.

In FIG. 1, the protected device 100 is connected by network equipment 110 to the Internet 120 or to another communication network, for example an Intranet or other local or wide-area communication network. The network equipment 110 may be a network router, a firewall with routing capability, a network dispatcher, a load balancer, or other equipment for supporting connections between the Internet 120 and the protected device 100. FIG. 1 also shows a correlation engine 140, which correlates network information from the network equipment 110 and intrusion information from the intrusion detection system 130 as explained below. The correlation engine 140 may be a programmable processor or its logical equivalent. In FIG. 1, the correlation engine 140 is shown apart from the other elements for purposes of clear discussion; nevertheless, the correlation engine 140 may be part of the intrusion detection system 130 as well

as a stand-alone element or otherwise integrated with other elements shown in FIG. 1. A network management center 150 may oversee the operation of the elements and system of FIG. 1.

When the intrusion detection system 130 detects an intrusion, it provides intrusion information 200 regarding the intrusion. FIG. 2 shows an exemplary structure of the intrusion information 200. In the exemplary structure of FIG. 2, the intrusion information 200 includes a source address 210, a destination address 220, a protocol type 230, a source port 240, and a destination port 250. The source address 210, destination address 220, protocol type 230, source port 240, and destination port 250 may be the respective elements of a message judged by the intrusion detection system 130 to be representative of messages that constitute the attack upon the protected device 100. For example, when the protected device 100 is a web server connected to the Internet 120, the protocol type 230 may be TCP, and the source address 210 and the destination addresses 220 of the message representative of the attack may be IP addresses. The port information, i.e., the source port 240 and the destination port 250, would concern the transport layer of the communication-protocol stack rather than logical ports on the network equipment 110. Those skilled in the art will appreciate, however, that the structure and particular elements of FIG. 2 are illustrative rather than limiting, and that the intrusion information 200 may come in different forms and may include different elements within the scope of the present invention.

When the intrusion detection system 130 detects an intrusion, the network equipment 110

provides related network information 300. FIG. 3 shows an exemplary structure of the network information 300. Although the network information 300 may take different forms depending on the nature of the network equipment 110, the exemplary network information 300 of FIG. 3 has the nature of a router table comprising three records 310 through 330, each of which describes a connection through the router. The records 310 through 330 may contain source addresses 310A through 330A, destination addresses 310B through 330B, protocol types 310C through 330C, logical input port identifiers 310D through 330D, and logical output port identifiers 310E through 330E, as shown in FIG. 3. The foregoing structure is illustrative, of course, rather than limiting, and other router implementations may store other information, for example MPLS labels, class of service information, and so forth.

In the intrusion information 200 and in the network information 300, the source addresses 210 and 310A through 330A and the destination addresses 220 and 310B through 330B need not necessarily be specific, single addresses; rather, they may also be ranges of addresses, or may be addresses that identify subnets, and so forth.

FIG. 4 shows an exemplary structure of a mapping table 400 for mapping the logical port identifiers 410A through 460A included in the network information 300 to physical port identifiers 410B through 460B of the network equipment 110. The table concerns network equipment 110 such as routers, and the mappings of FIG. 4 should not be confused with ports and mappings of the transport layer. The exemplary mapping table 400 of FIG. 4 shows six logical port identifiers 410A through 460A and six physical port identifiers 410B through 460B

in keeping with the need to map the six logical port identifiers 310D through 330D and 310E through 330E shown in the exemplary network information 300 of FIG. 3. Although FIG. 4 shows the mapping table 400 as a separate table, the mapping table 400 may be part of the network information 300 itself, or may be kept elsewhere; moreover, the mapping is not required to be one-to-one.

FIG. 5 shows aspects of the operation of the invention. The intrusion detection system 130 awaits an attack (step 500) upon the protected device 100. When an attack is not detected, the intrusion detection system 130 continues to wait (step 500). Otherwise (i.e., an attack is detected), the intrusion detection system 130 notifies the correlation engine 140 of the presence of an attack (step 505). The correlation engine 140 obtains intrusion information 200 from the intrusion detection system 130 (step 510) and network information 300 from the network equipment 110 (step 515).

The correlation engine 140 correlates the intrusion information 200 and the network information 300, looking for common elements (step 520). For example, the correlation engine 140 may search through the network information 300, looking for a record that has a source address that matches the source address 210 of the intrusion information 200. Alternatively, the correlation engine 140 may look through the network information 300 for a match of the destination address 220 of the intrusion information 200, or a match of the protocol type 230 of the intrusion information 200, or a match of two of the three elements 210 through 230 of the intrusion information 200, or a match of all three elements 210 through 230 of the intrusion information

200, or a match on other data which are not explicitly described herein but whose suitability would be evident to those skilled in the art once taught the present invention.

Finding a match as just described identifies one of the records 310 through 330 of the network information 300 that describes the connection through the network equipment 110 used by the attack. The analysis engine 140 examines the identified record to determine which of the logical input port identifiers 310D through 330D and the logical output port identifiers 310E through 330E is associated with the connection used by the attack (step 525). The correlation engine 140 then consults the mapping table 400 in order to map the logical input port and logical output port identified with the attack to the corresponding physical input port and physical output port of the network equipment 110 (step 530). The entry point of the attack into the protected device 100 has now been identified, and the correlation engine 140 alerts the network management center 150 (step 535) of the presence of attack and the entry point of the attack – i.e., the ports supporting the connection used by the attack – and then awaits another attack (step 500).

Although the invention is described above mainly in terms of finding the entry point of an attack upon a protected computer, the same methods may be applied as well to finding the exit point of an attack. Consequently, for descriptive convenience, the term “portal” is used herein inclusively; the portal of an attack may be an entry point of the attack or the exit point of the attack.

From the foregoing description, those skilled in the art will appreciate that the present invention improves the performance of equipment used to protect a computer, a web server, and so forth,

